

IN THE CLAIMS:

Please cancel Claims 1, 3-5, 10-15, 35-36, 43, and 54-55 without prejudice.

1 1. (Cancelled)

1 2. (Previously Presented) The method of Claim 6 wherein the step of associating
2 comprises the step of producing a result representing a remainder upon dividing the IP ID
3 by the number of active links.

1 3.-5. (Cancelled)

1 6. (Previously Presented) A method for uniformly distributing data transmitted by a
2 server over a plurality of underlying links of an aggregate within a computer network,
3 comprising:

4 defining a unit of data as a datagram;

5 apportioning each datagram into at least one fragment at the server;

6 associating each fragment to an underlying link of the aggregate on the basis of an
7 Internet protocol (IP) identifier (ID) of each datagram and a number of active links of the
8 aggregate, wherein the step of associating includes:

9 logically combining the IP ID with a predetermined mask to produce a quantity,

10 right shifting the quantity a predetermined number of places,

11 establishing a threshold at which a group of data is forwarded to each underlying
12 link of the aggregate,

13 producing a result representing a remainder upon dividing the right shifted logi-
14 cally combined quantity IP ID and predetermined mask by the number of active links,
15 wherein the IP ID is a 16-bit value, the predetermined mask is 0xFF80 and predetermined
16 number of right shifted places is 7, and wherein the group of data comprises 128 IP IDs;

17 transmitting the fragment over its associated underlying link from the server to the
18 computer network.

1 7. (Original) The method of Claim 6 wherein the group of data comprises one of 128
2 different transport control protocol (TCP) fragments and 128 different user datagram pro-
3 tocol (UDP) datagrams.

1 8. (Original) The method of Claim 7 wherein each UDP datagram comprises up to
2 23 fragments.

1 9. (Previously Presented) The method of Claim 1 further comprising:
2 loading at least one data buffer of the server with the at least one fragment;
3 fetching the fragment from the data buffer; and
4 loading at least one queue of the server with the fragment, the queue associated
5 with the underlying link.

1 10.-15.(Cancelled)

1 16. (Previously Presented) A computer readable medium, comprising:

2 the medium storing executable program instructions for uniformly distributing
3 data transmitted by a server over a plurality of underlying links of an aggregate within a
4 computer network, the executable program instructions having program instructions for:

5 defining a unit of data as a datagram;

6 apportioning each datagram into at least one fragment at the server;

7 associating each fragment to an underlying link of the aggregate on the basis of an
8 Internet protocol (IP) identifier (ID) of each datagram and a number of active links of the
9 aggregate, wherein the step of associating includes:

10 logically combining the IP ID with a predetermined mask to produce a quantity,

11 right shifting the quantity a predetermined number of places,

12 establishing a threshold at which a group of data is forwarded to each underlying
13 link of the aggregate,

14 producing a result representing a remainder upon dividing the right shifted logi-
15 cally combined quantity IP ID and predetermined mask by the number of active links,
16 wherein the IP ID is a 16-bit value, the predetermined mask is 0xFF80 and predetermined
17 number of right shifted places is 7, and wherein the group of data comprises 128 IP IDs;

18 transmitting the fragment over its associated underlying link from the server to the
19 computer network.

1 17. (Original) The computer readable medium of Claim 16 wherein the program in-
2 struction for associating comprises a program instruction for producing a result represent-
3 ing a remainder upon dividing the IP ID by the number of active links.

1 18. (Original) The computer readable medium of Claim 17 wherein the program in-
2 struction for associating further comprises program instructions for:

3 calculating the IP ID of each datagram in a sequential manner; and
4 rotating the fragments of each datagram among all the underlying links to thereby
5 ensure that all fragments having the same IP ID are provided to the same physical link of
6 the aggregate.

1 19. (Original) The computer readable medium of Claim 16 wherein the program in-
2 struction for associating comprises program instructions for:
3 logically combining the IP ID with a predetermined mask to produce a quantity;
4 right shifting the quantity a predetermined number of places; and
5 establishing a threshold at which a group of data is forwarded to each underlying link of
6 the aggregate.

1 20. (Original) The computer readable medium of Claim 19 wherein the program in-
2 struction for associating further comprises the program instruction for producing a result
3 representing a remainder upon dividing the right shifted logically combined quantity IP
4 ID and predetermined mask by the number of active links

1 21. (Previously Presented) A method for distributing data over a plurality of network
2 links within a computer network, comprising:
3 providing the plurality of links as a connection to a network node;
4 selecting one link of the plurality of links for transmitting a datagram to the net-
5 work node (hereinafter the selected link) using a round robin selection technique, the data
6 identified by an Internet protocol (IP) identifier (ID), the IP ID indicating an end point
7 destination for the data;
8 apportioning the datagram into at least one fragment;

9 associating the datagram to a network link of the plurality of network links ac-
10 cording to a round robin policy based at least in part on the IP ID;
11 associating each fragment to the selected link on the basis of the IP ID; and
12 transmitting the fragment over its associated link to transmit fragments of the data
13 over the same associated link

1 22. (Previously Presented) The method of claim 21 wherein the step of associating is
2 further based, at least in part, on a number of network links in the plurality of links.

1 23. (Previously Presented) The method of claim 22 wherein the step of associating
2 comprises:
3 producing a result representing a remainder by dividing the IP ID by the number
4 of network links.

1 24. (Previously Presented) The method of claim 23 wherein the step of associating
2 further comprises:
3 calculating the IP ID of each datagram in a sequential manner; and
4 rotating the fragments of each datagram among all the network links of the plural-
5 ity of network links to thereby ensure that all fragments having the same IP ID are pro-
6 vided to the same network link.

1 25. (Previously Presented) The method of claim 21 wherein the step of associating
2 comprises:
3 logically combining the IP ID with a predetermined mask to produce a quantity;

4 right shifting the quantity by a predetermined number of places to create a Previ-
5 ously Presented quantity; and

6 establishing a threshold at which a group of data is forwarded to each network
7 link of the plurality of network links.

1 26. (Previously Presented) The method of claim 25 wherein the step of associating
2 further comprises the step of producing a result representing a remainder upon dividing
3 the Previously Presented quantity by a number of network links in the plurality of net-
4 work links.

1 27. (Previously Presented) A system for distributing data over a plurality of network
2 links within a computer network comprising:

3 a processor;

4 a memory coupled to the processor and having locations addressable by the proc-
5 essor;

6 an operating system resident in the memory locations and executed by the proces-
7 sor, the operating system configured to implement a load balancing technique that defines
8 a unit of data as a datagram, the operating system further configured to include an Inter-
9 net Protocol (IP) layer that apportions the datagram into at least one fragment, the operat-
10 ing system further configured to include a virtual interface process that associates each
11 datagram to a network link of the plurality of network links according to a round robin
12 policy based at least in part on an Internet protocol (IP) identifier (ID) of each datagram;
13 and

14 at least one network adapter coupled to the memory and processor to transmit the
15 fragments of the datagram over the datagram's associated network link.

28. (Currently Amended) A system for distributing data over a plurality of network links within a computer network, comprising:

means for providing the plurality of links as a connection to a network node;

means for selecting one link of the plurality of links for transmitting a datagram to the network node (hereinafter the selected link) using a round robin selection technique, the data identified by an Internet protocol (IP) identifier (ID), the IP ID indicating an end point destination for the data;

means for defining a unit of data as a datagram;

means for apportioning the each datagram into a plurality of fragments at least one fragment;

means for associating each datagram to a network link of the plurality of network links according to a round robin policy based at least in part on an Internet protocol (IP) identifier (ID) of each datagram;

means for associating each fragment of the plurality of fragments to the selected link on the basis of the IP ID; and

means for transmitting the plurality of fragments over its associated link to transmit fragments of the data over the same associated link

means for transmitting the fragments of the datagram over the datagram's associated network link .

29. (Previously Presented) A computer readable medium, comprising:

the medium storing executable program instructions for execution on a processor,
the executable program instructions having program instructions for:

providing the plurality of links as a connection to a network node;
selecting one link of the plurality of links for transmitting a datagram to the network node (hereinafter the selected link) using a round robin selection technique, the data identified by an Internet protocol (IP) identifier (ID), the IP ID indicating an end point destination for the data;
apportioning the datagram into at least one fragment;
associating the datagram to a network link of the plurality of network links according to a round robin policy based at least in part on the IP ID;
associating each fragment to the selected link on the basis of the IP ID; and
transmitting the fragment over its associated link to transmit fragments of the data over the same associated link.

30. (Previously Presented) A method for distributing data over a plurality of network links within a computer network, comprising:

aggregating the plurality of network links to provide a connection to a network node, each link of the plurality of links assigned to the same MAC address of the network node;

dividing a first datagram, having a first Internet protocol (IP) identifier (ID), into one or more fragments, each fragment of the first datagram associated with the first Internet protocol (IP) identifier (ID);

selecting a first network link of the plurality of network links for transmission of the one or more fragments of the first datagram, the selection done by a round robin policy;

12 transmitting all of the one or more fragments associated with the first IP ID over
13 the first network link;

14 dividing a second datagram, having a second Internet protocol (IP) identifier (ID),
15 into one or more fragments, each fragment of the second datagram associated with the
16 second Internet protocol (IP) identifier (ID);

17 selecting a second network link of the plurality of network links for transmission
18 of the one or more fragments of the second datagram; and

19 transmitting all of the one or more fragments associated with the second IP ID
20 over the second network link.

1 31. (Previously Presented) The method of claim 30 wherein the first network link and
2 the second network link are selected according to a round robin policy based on the IP ID
3 of each datagram.

1 32. (Previously Presented) A system for distributing data over a plurality of network
2 links within a computer network comprising:

3 a processor;

4 a memory coupled to the processor and having locations accessible by the proces-
5 sor;

6 a virtual interface aggregating the plurality of network links to provide a connec-
7 tion to a network node, each link of the plurality of links assigned to the same MAC ad-
8 dress of the network node;

9 an operating system resident in the memory and executed by the processor, the
10 operating system configured to implement a load balancing technique that divides a first
11 datagram, having a first Internet protocol (IP) identifier (ID), into one or more fragments,
12 each fragment of the first datagram associated with the first Internet protocol (IP) identi-

13 fier (ID), select a first network link of the plurality of network links for transmission of
14 the one or more fragments of the first datagram, the selection done by round robin policy,
15 transmit all of the one or more fragments associated with the first IP ID over the first
16 network link, divide a second datagram, having a second Internet protocol (IP) identifier
17 (ID), into one or more fragments, each fragment of the second datagram associated with a
18 second Internet protocol (IP) identifier (ID), select a second network link of the plurality
19 of network links for transmission of the one or more fragments of the second datagram,
20 and transmit all of the one or more fragments associated with the second IP ID over the
21 second network link.

1 33. (Previously Presented) The system of claim 32 wherein the first network link and
2 the second network link are selected according to a round robin policy based on the IP ID
3 of each datagram.

1 34. (Previously Presented) The method of claim 6 wherein the step of associating fur-
2 ther comprises apportioning data equally over the plurality of underlying links of the ag-
3 gregate within the computer network.

1 35.-36.(Cancelled)

1 37. (Previously Presented) The system of claim 27 wherein the round robin policy is
2 further based, at least in part, on a number of network links in the plurality of links.

1 38. (Previously Presented) The system of claim 27 wherein the virtual interface proc-
2 ess is configured to produce a result representing a remainder by dividing the IP ID by
3 the number of network links.

1 39. (Previously Presented) The method of claim 38 wherein the virtual interface proc-
2 ess is further configured to calculate the IP ID of each datagram in a sequential manner
3 and to rotate the fragments of each datagram among all the network links of the plurality
4 of network links to thereby ensure that all fragments having the same IP ID are provided
5 to the same network link.

1 40. (Previously Presented) The system of claim 27 wherein the virtual interface proc-
2 ess is configured to logically combine the IP ID with a predetermined mask to produce a
3 quantity, right shift the quantity by a predetermined number of places to create a Previ-
4 ously Presented quantity, and establish a threshold at which a group of data is forwarded
5 to each network link of the plurality of network links.

1 41. (Previously Presented) The system of claim 40 wherein the virtual interface proc-
2 ess is further configured to produce a result representing a remainder upon dividing the
3 Previously Presented quantity by a number of network links in the plurality of network
4 links.

1 42. (Previously Presented) The method of claim 30, wherein the steps of transmitting
2 are performed in parallel.

1 43. (Cancelled)

1 44. (Currently Amended) A method for uniformly distributing data transmitted by a
2 server over a number of underlying links of an aggregate within a computer network, the
3 comprising:

4 providing the plurality of links as a connection to a network node;

5 selecting one link of the plurality of links for transmitting a datagram to the net-
6 work node (hereinafter the selected link) using a round robin selection technique, the data
7 identified by an Internet protocol (IP) identifier (ID), the IP ID indicating an end point
8 destination for the data;
9 apportioning the datagram into at least one fragment;
10 associating the fragments with the selected link; and
11 transmitting the fragments over the selected link.

1 45. (Previously Presented) The method of claim 44 wherein the step of calculating
2 further comprises:

3 dividing the IP ID by the number of underlying links to generate a remainder; and
4 using the remainder to select the selected link.

1 46. (Previously Presented) The method of claim 44 wherein the step of calculating
2 further comprises:

3 performing a logical AND operation to combine the IP ID and a predetermined
4 mask;

5 dividing the result of the logical AND operation by the number of underlying
6 links to generate a remainder; and

7 using the remainder as the link identifier.

1 47. (Previously Presented) The method of claim 46 wherein the predetermined mask
2 is 0xFF80.

1 48. (Previously Presented) The method of claim 44 wherein the step of calculating
2 further comprises:

3 performing a logical AND operation to combine the IP ID and a predetermined
4 mask;

5 right shifting the result of the logical AND by a predetermined number of bits;

6 dividing the result of right shifting by the number of underlying links to generate
7 a remainder; and

8 using the remainder as the link identifier.

1 49. (Previously Presented) The method of claim 48 wherein the predetermined mask
2 is 0xFF80 and the predetermined number of bits is 7 bits.

1 50. (Previously Presented) An apparatus for uniformly distributing data transmitted
2 by a server over a number of underlying links of an aggregate within a computer network,
3 comprising:

4 means for providing the plurality of links as a connection to a network node;

5 means for selecting one link of the plurality of links for transmitting a datagram to
6 the network node (hereinafter the selected link) using a round robin selection technique,
7 the data identified by an Internet protocol (IP) identifier (ID), the IP ID indicating an end
8 point destination for the data;

9 means for apportioning the datagram into at least one fragment;

10 means for associating the fragments with the selected link; and

11 means for transmitting the fragments over the selected link.

1 51. (Previously Presented) A computer readable medium storing executable program
2 instructions for uniformly distributing data transmitted by a server over a plurality of un-
3 derlying links of an aggregate within a computer network, the executable program in-
4 structions comprising program instructions for:

5 providing the plurality of links as a connection to a network node;
6 selecting one link of the plurality of links for transmitting a datagram to the net-
7 work node (hereinafter the selected link) using a round robin selection technique, the data
8 identified by an Internet protocol (IP) identifier (ID), the IP ID indicating an end point
9 destination for the data;
10 defining a unit of data as a datagram;
11 apportioning the datagram into at least one fragment;
12 associating the fragments with the selected link; and
13 transmitting the fragments over the selected link.

1 52. (Currently Amended) A method for uniformly distributing data transmitted by a
2 server over a number of underlying links of an aggregate within a computer network,
3 comprising:

4 providing the number of underlying links from the server as a connection to a
5 network node;
6 selecting one link of the plurality of links for transmitting data from the server to
7 the network node (hereinafter the selected link) using a round robin selection technique,
8 the data identified by an Internet protocol (IP) identifier (ID), the IP ID indicating an end
9 point destination for the data;
10 selecting a datagram for transmission, the datagram having ~~an~~ the Internet proto-
11 col (IP) identifier (ID);
12 breaking the datagram into fragments; and

13 sending all of the fragments down the selected link.

1 53. (Previously Presented) The method of claim 52, further comprising:
2 calculating a link identifier for the datagram as a function of the IP ID and the
3 number of underlying links;
4 associating the fragments with the link identifier; and
5 transmitting the fragments over a link identified by the link identifier.

1 54.-55. (Cancelled)

1 56. (Previously Presented) A method for distributing data over a plurality of network
2 links, comprising:
3 providing the plurality of links as a connection to a network node;
4 selecting one link of the plurality of links for transmitting the data to the network
5 node (hereinafter the selected link) using a round robin selection technique, the data identified by an Internet protocol (IP) identifier (ID), the IP ID indicating an end point destination for the data;
6 apportioning the data into at least one fragment;
7 associating each fragment to the selected link on the basis of the IP ID; and
8 transmitting the fragment over its associated link to transmit fragments of the data
9 over the same associated link.
10
11

1 57. (Previously Presented) The method of claim 56, further comprising:
2 obtaining a second data to transmit over the plurality of links; and

3 selecting by the round robin process a second link to transmit fragments of the
4 second data to the network node.

1 58. (Previously Presented) The method of claim 56, further comprising:
2 directing each link of the plurality of links to a common media access address
3 (MAC address).

1 59. (Previously Presented) A computer, comprising:
2 a plurality of links to provide a connection to a network node;
3 one link of the plurality of links selected for transmitting a data to the network
4 node (hereinafter the selected link) using a round robin selection technique, the data iden-
5 tified by an Internet protocol (IP) identifier (ID), the IP ID indicating an end point desti-
6 nation for the data;
7 a processor to apportion the data into at least one fragment;
8 software to associate each fragment to the selected link on the basis of the IP ID;
9 and
10 a port to transmit the fragment over its associated link to transmit fragments of the
11 data over the same associated link.

1 60. (Previously Presented) The computer as in claim 59, further comprising:
2 a second data, the second data to be transmitted over the plurality of links; and
3 the round robin process to select a second link to transmit fragments of the second
4 data to the network node.

1 61. (Previously Presented) The computer as in claim 59, further comprising:

2 a common media access address (MAC address) to which each link of the plural-
3 ity of links is directed.

1 62. (Previously Presented) A computer readable media, comprising:
2 said computer readable media containing instructions for execution on a processor
3 for a method of distributing data over a plurality of network links, the method having the
4 steps of:
5 providing the plurality of links as a connection to a network node;
6 selecting one link of the plurality of links for transmitting the data to the network
7 node (hereinafter the selected link) using a round robin selection technique, the data iden-
8 tified by an Internet protocol (IP) identifier (ID), the IP ID indicating an end point desti-
9 nation for the data;
10 apportioning the data into at least one fragment;
11 associating each fragment to the selected link on the basis of the IP ID; and
12 transmitting the fragment over its associated link to transmit fragments of the data
13 over the same associated link.